

Experimental Quantum Cryptography*

Charles H. Bennett
IBM Research[†]

François Bessette[‡]
Université de Montréal[§]

Gilles Brassard[¶]
Université de Montréal[§]

Louis Salvail[¶]
Université de Montréal[§]

John Smolin**
UCLA^{††}

September 1991

Abstract

We describe results from an apparatus and protocol designed to implement *quantum key distribution*, by which two users, who share no secret information initially: 1) exchange a random quantum transmission, consisting of very faint flashes of polarized light; 2) by subsequent public discussion of the sent and received versions of this transmission estimate the extent of eavesdropping that might have taken place on it, and finally 3) if this estimate is small enough, distill from the sent and received versions a smaller body of shared random information, which is certifiably secret in the sense that any third party's expected information on it is an exponentially small fraction of one bit. Because the system depends on the uncertainty principle of quantum physics, instead of usual mathematical assumptions such as the difficulty of factoring, it remains secure against an adversary with unlimited computing power.

* A preliminary version of this paper was presented at EUROCRYPT '90, May 21–24, Århus, Denmark, and has appeared in the proceedings, pp. 253–265.

[†] Yorktown Heights, New York, NY 10598, USA.

[‡] Supported in part by an NSERC Postgraduate Scholarship.

[§] Département IRO, Université de Montréal, C.P. 6128, succursale "A", Montréal (Québec), Canada H3C 3J7.

[¶] Supported in part by Canada's NSERC.

** This work was performed while this author was visiting IBM Research.

^{††} Physics Department, University of California at Los Angeles, Los Angeles, CA 90024, USA.

1 Introduction and History

Quantum cryptography has entered the experimental era [5]. The first convincingly successful quantum exchange took place in October 1989. After a short historical review of quantum cryptography, we report on the new apparatus and the latest results obtained with it.

Quantum cryptography was born in the late sixties when Stephen Wiesner wrote “Conjugate Coding”. Unfortunately, this highly innovative paper was unpublished at the time and it went mostly unnoticed. There, Wiesner explained how quantum physics could be used in principle to produce bank notes that would be impossible to counterfeit and how to implement what he called a “multiplexing channel”, a notion strikingly similar to what Rabin was to put forward more than ten years later under the name of “oblivious transfer” (in our opinion, it would be fair to give at least equal credit to Wiesner for the concept of oblivious transfer).

Fortunately, Charles H. Bennett knew Wiesner quite well and heard about his idea from the horse’s mouth. Nevertheless, it was only when he met Gilles Brassard that quantum cryptography was revived. This happened on the occasion of the 20th IEEE Symposium on the Foundations of Computer Science, held in Puerto Rico in October 1979. Following their discussion of Wiesner’s idea, they discovered how to incorporate the (almost new at the time) notion of public key cryptography, resulting in a CRYPTO ’82 paper [7]. This brought Wiesner’s paper back to life, and it was subsequently published in *Sigact News* [27], together with a selection of papers from the earlier CRYPTO ’81 workshop.

Initially, quantum cryptography was thought of by everyone (including ourselves) mostly as a work of science-fiction because the technology required to implement it was out of reach (for instance, quantum bank notes [27] require the ability to store a single polarized photon or spin- $\frac{1}{2}$ particle for days without significant absorption or loss of polarization). Unfortunately, the impact of the CRYPTO ’82 conference had left most people under the impression that everything having to do with quantum cryptography was doomed from the start to being unrealistic.

The main breakthrough came when Bennett and Brassard realized that photons were never meant to *store* information, but rather to *transmit* it (although it should be said that half of Wiesner’s original paper dealt precisely with the use of quantum physics for the transmission of information). This lead initially to the *self-winding reusable one-time pad* [6] which was still not very practical. Later, Bennett thought of the *quantum key distribution channel* (whose implementation is the object of the current paper) and Brassard designed the somewhat less realistic *quantum coin-tossing protocol* (which can be used to implement bit commitment) [2, 3]. Quantum cryptography was also picked up by other researchers. For instance, Crépeau and Kilian showed how the quantum channel could be used in principle (although not in practice) to implement oblivious transfer in a strong way (Wiesner’s original multiplexing channel could leak information on both channels), zero-knowledge protocols, and secure two-party computation [17, 16]. More recently, Ekert proposed an alternative

approach to implement quantum key distribution [19] (making use of EPR and Bell’s theorem), but a simplified —and no less secure— version of his scheme is shown in [10] to be equivalent to the idealized quantum key distribution protocol originally put forward by Bennett and Brassard in 1984 [3]. Let us also mention that Bennett, Brassard, and Crépeau have developed practical quantum protocols to achieve oblivious transfer, bit commitment and coin-tossing [8]. See also [14].

The principle of quantum cryptography has been described in major popular magazines such as *Scientific American* [25], *The Economist* [20], *New Scientist* [18] and *Science News* [23]. In *New Scientist*, Deutsch wrote that “Alan Turing’s theoretical model is the basis of all computers. Now, for the first time, its capabilities have been exceeded” [by the quantum cryptography apparatus]¹ [18]. Also, Brickell and Odlyzko close their thorough survey of recent (1988) results in cryptanalysis with these words: “If such systems [quantum cryptography] become feasible, the cryptanalytic tools discussed here [in their paper] will be of no use” [15].

In this paper, we report on the first experimental quantum key distribution channel ever designed and actually put together. Section 2 provides background information on quantum cryptography. (For further detail on the basic quantum channel, see chapter 6 of [13].) We first review the original quantum key distribution protocol of [3], which illustrates the method most plainly. Then, we describe subsequent modifications of the protocol [11, 12, 4], which give it the ability, necessary in practice, to function despite partial information leakage to the eavesdropper and partial corruption of the quantum transmissions by noise. In Section 3, we describe the physical apparatus by which quantum key distribution has actually been carried out. In Section 4, we discuss the possible sources of information leakage to the eavesdropper. In Section 5, we report on actual data transmitted by the apparatus. Finally, the Appendix gives a new technique allowing privacy amplification to be applied when the eavesdropper’s information is probabilistic.

2 Quantum Key Distribution

The purpose of key distribution is for two users “Alice” and “Bob”, who share no secret information initially, to agree on a random key, which remains secret from an adversary “Eve”, who eavesdrops on their communications. In conventional cryptography and information theory it is taken for granted that digital communications can always be passively monitored, so that the eavesdropper learns their entire contents, without the sender or receiver being aware that any eavesdropping has taken place. By contrast, when digital information is encoded in elementary quantum systems such

¹ More precisely, it is mathematically impossible for two probabilistic interactive Turing machines who share only a short secret key beforehand to achieve secure exchange of a longer secret key under the nose of a third Turing machine eavesdropping on all their communications if that third machine has unlimited computing power. In sharp contrast, this is precisely what the experimental quantum cryptography prototype achieves, with an arbitrarily small probability of failure. This does not contradict the Church–Turing thesis since the purpose of the apparatus is *not* to compute functions.

as single photons, it becomes possible to produce a communications channel whose transmissions cannot in principle be reliably read or copied by an eavesdropper ignorant of certain information used in forming the transmission. The eavesdropper cannot even gain partial information about such a transmission without disturbing it in a random and uncontrollable way likely to be detected by the channel's legitimate users.

The essential quantum property involved, a manifestation of Heisenberg's uncertainty principle, is the existence of pairs of properties that are incompatible in the sense that measuring one property necessarily randomizes the value of the other. For example, measuring a single photon's linear polarization randomizes its circular polarization, and vice versa. More generally any pair of polarization states will be referred to as a *basis* if they correspond to a reliably measurable property of a single photon, and two bases will be said to be *conjugate* [27] if quantum mechanics decrees that measuring one property completely randomizes the other. Our quantum key distribution protocol uses two conjugate bases, which we shall take to be the rectilinear basis (horizontal vs vertical polarization) and the circular basis (left-circular vs right-circular). We shall refer to these as the *canonical bases*. Similarly, a *canonical polarization* is either horizontal, vertical, left-circular, or right-circular. A third basis also exists, consisting of 45 and 135 degree diagonal polarizations, which is conjugate to both the other two bases, but we will not need to consider it except in connection with possible eavesdropping strategies. More information on the notion of conjugate bases is given in the Appendix.

The protocol we describe here is secure even against an enemy possessing unlimited computing power (even if $\mathcal{P} = \mathcal{NP}$!), under any attack in which she is limited to measuring photons (or in the subsequent generalization, light pulses) one at a time, and combining the classical results of these measurements with information subsequently overheard during the public discussion (described below). The formalism of quantum mechanics allows a more general kind of measurement, completely infeasible at present or in the foreseeable future. Such a measurement would treat the entire sequence of n photons sent during a key-distribution session as a single 2^n -state quantum system, cause it to interact coherently with an intermediate quantum system of comparable complexity, maintain the phase coherence of the intermediate system for an arbitrarily long time, then finally measure the intermediate system in a way depending on the information overheard during the public discussion. It is not known whether the protocol is secure against such an attack, but recent work indicates that it may be [10].

The basic quantum key distribution protocol (cf Figure 1) begins with Alice sending a random sequence of the four canonical kinds of polarized photons to Bob. Bob then chooses randomly and independently for each photon (and independently of the choices made by Alice, of course, since these choices are unknown to him at this point) whether to measure the photon's rectilinear or circular polarization. Bob then announces publicly which kind of measurement he made (but not the result of the measurement), and Alice tells him, again publicly, whether he made the correct

1.	\curvearrowright	\updownarrow	\curvearrowleft	\leftrightarrow	\updownarrow	\updownarrow	\leftrightarrow	\leftrightarrow	\curvearrowleft	\curvearrowright	\updownarrow	\curvearrowleft	\curvearrowright	\updownarrow
2.	+	○	○	+	+	○	○	+	○	+	○	○	○	+
3.	\updownarrow		\curvearrowleft		\updownarrow	\curvearrowright	\curvearrowright	\leftrightarrow		\updownarrow	\curvearrowleft	\curvearrowleft		\curvearrowright
4.	+		○	+	○	○	+		+	○	○		○	+
5.			✓		✓			✓			✓		✓	✓
6.			\curvearrowleft		\updownarrow		\leftrightarrow			\curvearrowleft		\curvearrowright	\updownarrow	
7.			1		1		0			1		0	1	

Figure 1: Basic quantum key distribution protocol.

1. Alice sends a random sequence of photons polarized horizontal (\leftrightarrow), vertical (\updownarrow), right-circular (\curvearrowright) and left-circular (\curvearrowleft);
2. Bob measures the photons' polarization in a random sequence of bases, rectilinear (+) and circular (○).
3. Results of Bob's measurements (some photons may not be received at all).
4. Bob tells Alice which basis he used for each photon he received;
5. Alice tells him which bases were correct;
6. Alice and Bob keep only the data from these correctly-measured photons, discarding all the rest.
7. This data is interpreted as a binary sequence according to the coding scheme $\leftrightarrow = \curvearrowright = 0$ and $\updownarrow = \curvearrowleft = 1$.

measurement (ie rectilinear or circular). Alice and Bob then agree publicly to discard all bit positions for which Bob performed the wrong measurement. Similarly, they agree to discard bit positions where Bob's detectors failed to detect the photon at all — a fairly common event with existing detectors at optical wavelengths. The polarizations of the remaining photons is interpreted as bit 0 for horizontal or left-circular, and bit 1 for vertical and right-circular. The resulting binary string should be shared secret information between Alice and Bob, provided that no eavesdropping on the quantum channel has taken place. The result of the above steps is referred to as the *quantum transmission* (or sometimes the *raw* quantum transmission to emphasize that it was obtained early in the process).

In the basic protocol, Alice and Bob next test for eavesdropping by publicly comparing polarizations of a random subset of the photons on which they think they should agree. As shown in the Appendix, no measurement the eavesdropper can make on one of these photons while it is in transit from Alice to Bob can yield more than $\frac{1}{2}$ expected bit of information on its polarization. Moreover, any measurement yielding $s \leq \frac{1}{2}$ expected bit has probability at least $s/2$ of inducing a discrepancy when the data of Bob and Alice are compared, assuming that this photon is detected in the correct basis by Bob (otherwise, this photon is lost to all parties). If Alice and

Bob find no discrepancies, and if it is safe to assume that Eve cannot corrupt the contents of the public messages exchanged between them, then Alice and Bob may safely conclude that there are few or no errors in the remaining uncomparated data, and that little or none of it is known to any eavesdropper.

The assumption that the public messages cannot be corrupted by Eve is necessary, because otherwise it is clear that Eve could sit between Alice and Bob and impersonate each of them to the other [1]. As a result, Eve would end up with a string shared with Alice and another one shared with Bob, whereas Alice and Bob would be none the wiser. This crucial property of the public channel can be implemented in practice either by using an inherently unjammable public channel or by using an information-theoretically secure authentication scheme [26] to certify that the public messages have not been altered in transit. In the latter case, Alice and Bob need to have a modest amount of shared secret information beforehand to serve as an authentication key, and a few bits of this key are rendered unfit for re-use each time the key distribution protocol is carried out. However, each successful instance of the protocol provides Alice and Bob with a substantially larger volume of fresh key information, some of which can be used to replace the lost authentication bits. Hence in this case the protocol implements key *expansion* rather than key distribution. It should be noted that in the case of a jammable public channel, a determined opponent, by repeated interference with either the quantum or public transmissions, could force Alice and Bob to exhaust their entire supply of authentication key before successfully distributing any fresh key to replace it. Alice and Bob would then have irreversibly lost their ability to exchange key securely. Nevertheless, except with arbitrarily small probability, Eve cannot make Alice and Bob believe that they have succeeded when in fact their fresh key information is either not shared or not secret, or both.

The elementary “quality-control” in the basic quantum key distribution protocol, which follows the quantum transmission as described above, is inadequate in practice for two reasons:

1. Realistic detectors have some noise; therefore, Alice’s and Bob’s data will differ even in the absence of eavesdropping. Accordingly, they must be able to recover from a reasonably small error frequency.
2. It is technically difficult to produce a light pulse containing exactly one photon. It is much easier to produce a coherent pulse, which may be regarded as a superposition of quantum states with 0, 1, 2... photons; or an incoherent pulse, which may be regarded as a statistical mixture of coherent states. In either case, let μ be the expected number of photons per pulse. If μ is small (ie significantly less than 1), there is a probability approximately $\mu^2/2$ that an eavesdropper will be able to split a pulse into two or more photons, reading one and allowing the other(s) to go to Bob². This allows the eavesdropper to learn a constant fraction of the bits shared between Alice and Bob without inducing errors.

² More generally, the probability of detecting k photons in a single pulse is given (exactly for coherent light, and in the low-intensity limit for incoherent light) by a Poisson distribution of mean μ .

A satisfactory protocol must be able to recover from noise as well as from partial leakage.

Below we describe a practical protocol that remedies these defects, allowing Alice and Bob to reconcile the differences between the sent and received versions of the quantum transmission, and then distill from the reconciled data (about which the eavesdropper may have significant partial information) a smaller body of data that is almost perfectly secret. The protocol we sketch is simple but not optimal: other protocols, which we are currently developing, have a higher yield of shared secret key at the same levels of noise and leakage. Further details on preliminary versions of the current protocol may be found in [4, 11, 24, 12].

Once the quantum transmission has been completed (with very dim light pulses used instead of single photons, as discussed in (2) above), the first task is for Alice and Bob to exchange public messages enabling them to reconcile the differences between their data. Because we assume throughout that Eve listens to all the public messages between Bob and Alice, this exchange must be performed in a way that reveals as little information as possible on this data. On the other hand, let us recall that Eve cannot corrupt the contents of these public messages.

An effective way for Alice and Bob to perform reconciliation is for them first to agree on a random permutation of the bit positions in their strings (to randomize the locations of errors), then partition the permuted strings into blocks of size k such that single blocks are believed to be unlikely to contain more than one error. (The optimal block size, which should be a function of the expected error rate, has not yet been determined theoretically; instead, we use in Section 5 block sizes that have empirically been found to be good.) For each such block, Alice and Bob compare the block's parity. Blocks with matching parity are tentatively accepted as correct, while those of discordant parity are subject to a bisection search, disclosing $\log(k)$ further parities of sub-blocks, until the error is found and corrected. If the initial block size was much too large or too small, due to a bad a priori guess of the error rate, that fact will become apparent, and the procedure can be repeated with a more suitable block size³. In order to avoid leaking information to Eve during the reconciliation process, Alice and Bob agree to discard the last bit of each block or sub-block whose parity they have disclosed.

Of course, even with an appropriate block size, some errors will typically remain undetected, having occurred in blocks or sub-blocks with an even number of errors. To remove additional errors, the random permutation and block parity disclosure is repeated several more times, with increasing block sizes, until Alice and Bob estimate that at most a few errors remain in the data as a whole. At this point, the block parity disclosure approach becomes much less efficient because it forces Alice and Bob to sacrifice at least one bit in each block on the altar of privacy. Consider for instance a (very typical) situation in which exactly two errors are left. If the block size is

³ Alternatively, a small random sample of the bits could be compared initially in order to estimate the error rate, much like the quality control mechanism in the basic quantum key distribution protocol. Of course, these bits would then have to be sacrificed.

chosen so that there are ℓ blocks, the probability of *not* detecting the existence of the remaining errors is $1/\ell$, and the cost for this strategy is ℓ bits when unsuccessful. For this reason, a different strategy is adopted to eliminate any errors that may remain and to verify, with high probability, that they have in fact been eliminated. The probability of undetected errors with this new strategy is $2^{-\ell}$ for the same cost of ℓ bits sacrificed to privacy, and this probability is completely independent of the number and location of remaining errors.

In each iteration of this strategy, Alice and Bob compare parities of a publicly chosen random subset of the bit positions in their entire respective data strings. If the data strings are not identical, then the random-subset parities will disagree with probability exactly $\frac{1}{2}$. If a disagreement is found, Alice and Bob undertake a bisection search, similar to that mentioned above, to find and remove the error. As in the preceding block-parity stage of the reconciliation, the last bit of each compared subset is discarded to avoid leaking any information to Eve. Each subsequent random subset parity is, of course, computed with a new independent random subset of bit positions in the remaining string.

At some point, all errors will have been removed, but Alice and Bob will not yet be aware of their success. When this occurs, subsequent random subset parities will of course always agree. After the last detected error, Alice and Bob continue comparing random subset parities until sufficiently many consecutive agreements (say 20) have been found to assure them that their strings are indeed identical, with a negligible probability of not detecting the existence of remaining errors.

Alice and Bob are now in the possession of a string that is almost certainly shared, but only partly secret. As described in Section 4, they can find a conservative estimate on Eve’s partial information on their string from the detected error frequency and the optical pulse intensity. More precisely, they can estimate an integer ℓ such that Eve’s information on Alice’s string resulting from the raw quantum transmission is worth no more than knowledge of ℓ physical bits of that string (please consult the Appendix for the meaning of “worth no more”). Recall that the reconciliation process involves Alice disclosing the parity of many subsets of her bits, but that each time one bit from that subset is discarded from the reconciled string. As a result, Eve’s knowledge about physical bits could become knowledge about parities. Let us say that Eve *knows a parity bit* about Alice’s string if she knows the parity of a non-empty subset of the bits of that string (knowledge of physical bits is a special case of knowledge of parity bits, taking single-element subsets). It is easy to see that if Eve knows no more than ℓ parity bits about a string y , and if she is given an additional parity bit about y , but that z is formed by discarding from y one of the bits involved in that parity, then Eve still knows no more than ℓ parity bits about z . Therefore, if Eve knew no more than ℓ physical bits of Alice’s string before reconciliation, she knows no more than ℓ parity bits about the string shared between Alice and Bob that results from reconciliation.

At this point, Alice and Bob can perform *privacy amplification*, which is a fundamental tool introduced in [12]. Let x denote the reconciled string and let n denote

its length. Let us say that a *deterministic bit* of information about x is the value $e(x)$ of an arbitrary function $e : \{0, 1\}^n \rightarrow \{0, 1\}$. For instance, physical and parity bits are deterministic bits, but bits of information in the sense of Shannon's information theory need not be. It is shown in [12] that if Eve's knowledge about x is no more than ℓ deterministic bits, a hash function h randomly and publicly chosen from an appropriate class of functions $\{0, 1\}^n \rightarrow \{0, 1\}^{n-\ell-s}$ will map x into a value $h(x)$ about which Eve's expected information is less than $2^{-s}/\ln 2$ bit, where $s > 0$ is an arbitrary security parameter. This technique applies for Alice and Bob because parity bits are a special case of deterministic bits.

An adequate hash function for this purpose can be obtained by continuing to compute $n - \ell - s$ additional publicly chosen independent random subset parities, but now keeping their values secret instead of comparing them. The class of hash functions thus realized is essentially the strongly-universal₂ class H3 discussed by Wegman and Carter [26]. It is amusing to note that if even a single discrepancy is left between Alice's and Bob's data after reconciliation, the final strings computed by Alice and Bob will be totally uncorrelated, a fact likely to be noticed rapidly. Moreover, it is clear that this hash function has the property that if Eve's knowledge of x before privacy amplification was strictly in the form of parity bits, then such is also the case about her knowledge of $h(x)$. Therefore, Eve cannot have nonzero information about $h(x)$ without in fact having at least one bit of information about it. As a consequence, the privacy amplification theorem implies that Eve knows *nothing at all* about the final string $h(x)$ shared between Alice and Bob, except with probability at most $2^{-s}/\ln 2$, in which case she knows at least one deterministic bit.

3 Physical Apparatus

The apparatus (cf Figure 2) occupies an optical bench approximately one meter long inside a light-tight box measuring approximately $1.5 \times .5 \times .5$ meters. It is controlled by a program running on an IBM PC computer, which contains separate software representations of the sender Alice, who controls the sending apparatus, the receiver Bob, who controls the receiving apparatus, and optionally an eavesdropper Eve. The program can also run in simulation mode, without the attached experimental apparatus. Even though they reside in the same computer, no direct communication is allowed between the software Alice and the software Bob, except the public channel communication called for by the protocol.

Alice's light source, at the left end of the optical bench, consists of a green light-emitting diode (LED Stanley type HBG5566X) as the source of incoherent light, a 25 micron pinhole and 25 mm focal length lens to form a collimated beam, a 550 ± 20 nm interference filter (Ealing type 35-5065) to reduce the intensity and spectral width of the light and select a portion of the spectrum at which the photomultipliers have relatively high quantum efficiency, and finally a Polaroid filter (ie a dichroic sheet polarizer) to polarize the beam horizontally. The LED is driven by current pulses

Figure 2: Photograph of the apparatus.

Incoherent green light flashes are produced by Light Emitting Diode (LED) on the left, collimated into a beam by Pinhole and Lens, then pass through a 550 nm Filter and a horizontal Polarizer. Sender's Pockels Cells convert the horizontal polarization into an arbitrary sequence of the four polarization states (horizontal, vertical, left-circular and right-circular). After traversing the quantum channel, a 32 cm free air optical path, the beam passes through Receiver's Pockels Cell, which, if energized, converts rectilinear into circular polarizations and vice versa. Finally, a calcite Wollaston prism splits the beam into horizontally and vertically polarized components, in which individual photons are detected by Photomultiplier tubes A and B, respectively.

(about 5×10^{-8} coulombs in 60 nanoseconds) yielding, after collimation, filtration and polarization, an intensity of about 0.1 photon per pulse, about half of which is emitted during the first 500 nanoseconds. The low intensity of the light pulse serves to minimize the chance that an eavesdropper will be able to split any one pulse into two or more photons.

Alice modulates the polarization of the beam by means of two Pockels cells (INRAD type 102-020), operated at + or - the quarter-wave voltage (about 800 volts), so as to be able to choose among the four polarization states {horizontal, vertical, left-circular, or right-circular}. (Diagonal polarizations could have been used instead of circular, but they would have required twice the Pockels cell voltage.) High voltage NPN transistors (type BU-205), in series with 200K ohm pull-up resistors, are used to switch the high voltage for the Pockels cells under control of low voltage TTL signals on output lines of the PC's parallel port (5.1 volt Zener diodes protect the computer from exposure to high voltage in case of transistor failure).

The quantum channel itself is a free air optical path of approximately 32 centimeters.

Bob's receiving apparatus, at the right end of the optical bench, consists of another Pockels cell and a calcite Wollaston prism (Melles-Griot type 03PPW001/C), oriented so as to split the beam into vertically and horizontally polarized beams, which are directed into two photomultiplier tubes (Hamamatsu type R1463-01) with integral

preamplifiers and voltage dividers in the sockets (Hamamatsu type C716-05). Bob's Pockels cell is also operated at quarter wave voltage, allowing him to use the same Wollaston prism to make a measurement of either rectilinear or circular polarization, depending on whether the voltage is off or on.

The timing for each experiment is controlled by a timing and detection unit, which also contains the hardware for handling asynchronous communication with the PC's parallel port, and two potentiometers for setting the discrimination levels for rejecting small pulses from each photomultiplier preamplifier (no rejection of large pulses is necessary, owing to their infrequency). The pulse-height discrimination is carried out by fast ECL voltage comparators (Plessey type SP9687).

Upon receiving a "start" signal on one of the PC parallel port's output lines, the timing unit waits 60 μ sec for the Pockels cell voltages to settle, supplies current to the LED for about 60 nsec, gates the photomultiplier detection logic on for about 500 nsec, and sets two input lines of the parallel port according to the result (for each photomultiplier, whether a count was detected). The 500 nsec time window was chosen to include the brightest part of each light pulse (about half the integrated intensity), while avoiding excessive dark counts that would have accumulated had the window been kept open for the entire ≈ 5 microsecond duration of the pulse. When it has done all this, the timing unit turns on another of the parallel port's input lines to signify "done", and begins waiting for the next start signal. When the computer sees the done signal it knows it can read the results of the present experiment and thereafter safely start the next experiment.

Alice's choice of polarization and Bob's choice of reading basis are made randomly (not pseudorandomly) using a large file of random bits supplied to the computer on a diskette. Of course, Alice and Bob feed on different bits from this diskette (recall that although they live on the same computer, they do not communicate or otherwise share information that is not called for by the public channel discussion). These random bits had been previously generated using the same experimental apparatus, by taking the physically random output of one of the photomultipliers, illuminated by an auxiliary nearby LED of intensity such as to yield a count in about $\frac{1}{2}$ the time windows, removing the 0/1 bias by von Neumann's trick (in each consecutive pair of coin tosses, take HT=1, TH=0, and ignore HH and TT), and xORing the resulting bits with pseudorandom bits from the computer to hide any residual deviations from randomness caused by time-variation of the photomultiplier and pulse-detection circuit. The same file is used to supply additional random bits as needed by Alice and Bob during the data reconciliation and privacy amplification protocols described in the previous section.

The photomultipliers have quantum efficiency approximately 9%, with dark count rates of about 200 per second, or about 10^{-4} per 500 nsec time window. When using pulses of 0.1 expected photon per pulse, with about 0.05 expected photon arriving within the 500 nsec time window, this dark count rate would yield a bit error rate around 2%; the actual error rate, about 4%, was due also to imperfect alignment of the Pockels cells.

The driver program on the PC provides the ability to simulate two principal kinds of eavesdropping: intercept/resend and beamsplitting (described in Section 4) by a hypothetical adversary “Eve” who has detectors of 100% quantum efficiency.

The present apparatus is only an experimental prototype. In a more realistic demonstration, the error rate could be reduced several orders of magnitude by better optical alignment and cooling the photomultipliers to reduce dark current, the quantum channel could be made much longer, and the protagonists Alice, Bob, and Eve could reside in separate buildings [4]. The feasible distance over which a quantum key distribution system can operate depends on the noise and quantum efficiency of the detectors and especially on the attenuation of the optical channel: the weak signal entering the channel must still be recognizable above background upon leaving the channel. This means that essentially unlimited distances could be realized in principle by sending a diffraction-limited beam through an evacuated pipe or periscope of appropriate dimensions (eg 1 meter \times 1000 kilometers), using high-reflectance polarization preserving mirrors to change the beam direction as necessary. More practically, existing optical fibers have sufficiently low attenuation to be used over a distance of at least several kilometers.

It should be stressed that although the current prototype is of no direct practical value because no one is interested in secure key distribution over short distances, other cryptographic primitives make perfect sense in this context. For instance, an implementation of quantum oblivious transfer over a short distance is potentially very useful [8], as it would among other things allow a practical and unbreakable implementation of secure two-party computation.

4 Eavesdropping Strategies

This section describes the main eavesdropping strategies —intercept/resend and beamsplitting— and how they can be simulated with our apparatus and software. It also explains how Alice and Bob can estimate the amount of information potentially leaked to Eve. We assume conservatively throughout that Eve has unlimited technology (consistent with quantum physics) for dealing with single light pulses, including perfect photodetection and the capability of storing a pulse for an arbitrary long time before measuring it. However, we do not grant Eve the technology necessary to perform coherent measurements, as described in the third paragraph of Section 2.

4.1 Intercept/Resend

Recall that μ is the expected number of photons per light pulse. If μ is sufficiently smaller than 1, it is approximately also the probability that a pulse would be detected by a perfectly efficient detector. Let us say that a pulse is *successful* if Bob detects it in the basis originally chosen by Alice. In other words, successful pulses are those that contribute a bit in the raw quantum transmission. Unsuccessful pulses can be

ignored, except for the fact that Alice and Bob would become suspicious should the rate of successful pulses be significantly different from expected. (With the present apparatus about one pulse in 400 is successful).

In intercept/resend, Eve intercepts selected light pulses and reads them in bases of her choosing. For each such pulse, with probability approximately μ , Eve's perfectly efficient detectors are successful at detecting a photon. When this occurs, Eve fabricates and sends to Bob a pulse of the same polarization as she detected. To avoid suspicion with respect to the rate of successful pulses, Eve's fabricated pulses should be of such intensity (slightly higher than one expected photon per pulse) as to yield the same net rate of pulse detection by Bob as if no eavesdropping were taking place.

In a classical communications channel, it would be possible for Eve to measure Alice's signal exactly, and resend an exact copy of it, thereby escaping detection. However, in the present quantum setting, it is shown in the Appendix that at least 25% of the pulses Eve fabricates will yield the wrong result if later successfully measured by Bob. Moreover, each of these intercept/resends is worth no more to Eve than if she were told Alice's bit with probability $1/\sqrt{2}$, and told nothing with complementary probability.

Thus if there are t errors in the raw quantum transmission, Alice and Bob may conservatively estimate that fewer than $4t + 5\sqrt{12t}$ of their bits have been subjected to intercept/resend, the second term in this expression being an arbitrary but generous 5 standard-deviation allowance for statistical error. Furthermore, Alice and Bob may also conservatively estimate that the amount of information leaked to Eve through intercept/resend is worth no more than if she had obtained $(4/\sqrt{2})t + 5\sqrt{(4 + 2\sqrt{2})t}$ bits of the quantum transmission (the standard deviation of this estimator is explained in the Appendix). Of course, errors in practice may also result from other causes, such as optical misalignment, disturbance on the quantum channel, or noise in Bob's detectors. Therefore, Alice and Bob might be tempted to determine empirically the expected error rate in absence of eavesdropping and use the difference between the observed and predicted error rates to estimate the amount of information leaked to Eve through intercept/resend. However, this approach could grossly underestimate Eve's capabilities. Consider for instance the case in which most of the errors are due to misalignment between Alice's and Bob's notion of "horizontal". A clever Eve might discover this and rotate Alice's signal (eg by passing it through a sugar solution of the appropriate concentration) just enough to correct the misalignment and so reduce the error rate dramatically. Having done this, Eve could then intercept/resend a number of photons sufficient to create as many errors as the misalignment would have caused had Eve not corrected it. Clearly, such eavesdropping would completely avoid detection. In order to thwart this threat, it is safer if Alice and Bob take the very conservative view that *all* errors are due to intercept/resend.

No additional hardware is needed to simulate intercept/resend: when the software Eve wishes to intercept a pulse, she borrows the real receiving apparatus from Bob; when she wishes to resend to Bob, she borrows the sending apparatus from Alice. While Eve is borrowing the receiving apparatus, Alice obliges her by repeating the

same transmission $1/q$ times, where q is the quantum efficiency of the actual detectors. This allows Eve to obtain a count with the same probability μ as a physical eavesdropper with perfectly efficient detectors. Similarly, while Eve is borrowing the sending apparatus, Bob obliges her by allowing her to repeat the transmission $1/\mu$ times.

4.2 Beamsplitting

The other attack, beamsplitting, depends on the fact that the transmitted light pulses are not pure single-photon states. To carry out this attack, Eve uses a partly-silvered mirror or equivalent device to divert a fraction f of the original beam's intensity to herself, letting the remainder pass undisturbed to Bob. In order to avoid wasting information by measuring pulses in the wrong bases, Eve stores her share of each pulse until the correct bases have been announced in the public discussion. Then, Eve measures her stored pulses in those bases. With probability approximately $f\mu$, Eve will succeed in detecting a photon, and will therefore obtain Alice's bit for that pulse. This attack induces no errors, but does reduce the intensity reaching Bob by a factor $1 - f$.

A small loss of intensity might not be detectable to Alice and Bob, and in the present setup Eve could supplement the fraction f she splits from the main pulse by diverting to herself all the remaining integrated beam intensity during the 5 microsecond tail (cf Section 3), which Bob shuts out to avoid excessive dark counts, resulting in an effective splitting ratio of about $(f + 1)/2$. Another approach by which Eve can bring her share f closer to 1 without attenuating the signal too substantially would be for her to resend bright pulses when she intercept/resends pulses as described in Subsection 4.1. In analyzing data from the present apparatus, we conservatively take $f = 1$ and assume that Eve learns a fraction μ of Alice's bits through beamsplitting.

More realistically, if Eve does not have the technology needed to store her share of the beam, she can measure her share of each pulse as it comes. Of course, this can only be done at the cost of losing information. If Eve's measurements are in the canonical bases, she will learn a fraction roughly $\mu/2$ of Alice's string (still conservatively taking $f = 1$). If she uses noncanonical bases, the technique introduced in the Appendix can be used to show that she cannot be in a better situation than if she had obtained a fraction $\mu/\sqrt{2}$ of Alice's string. It should be noted that even if Eve had the capability of storing beams, but if Alice and Bob suspected it, they could send and receive all the pulses first, wait an arbitrarily long time for Eve's stored beam to decay, and only then announce all the bases.

The open air optical path in the present apparatus has negligible attenuation, but quantum key distribution might also be attempted through a channel with considerable natural attenuation, such as an optical fiber. If in addition Eve has the technological ability to surreptitiously substitute a much more transparent channel, she can mount an aggressive version of the beamsplitting attack based on allowing Bob to receive only pulses she has already succeeded in splitting. To do this, Eve

splits all the pulses entering the channel into two pulses of half intensity $\mu/2$, attempts to measure one half-pulse, and, if she succeeds in detecting a photon, forwards the other half-pulse to Bob, otherwise stopping it. She compensates for the resulting $(2/\mu)$ -fold reduction in pulse frequency and 2-fold reduction in intensity by increasing the channel's transparency $(4/\mu)$ -fold, so Bob receives photons at the same rate he would have from an undisturbed transmission through the original less transparent channel. This attack can be thwarted by keeping the original pulse intensity μ small compared to $4T$, where T is the transmission coefficient of the original channel before possible improvement by Eve.

A dramatic but harmless variant on the above attack would be for Eve to attempt to detect enough photons in the incoming pulse to determine its polarization uniquely, even without knowing the correct basis. An example of such a measurement would be for Eve to split the incoming pulse into two half-pulses as before, but now measure the rectilinear polarization of one and the circular polarization of the other. If, by extreme good luck, this measurement yielded three photons with polarizations vertical, horizontal, and right-circular, Eve would know that the original pulse's polarization was definitely right-circular, and she could capitalize on this knowledge by sending Bob such a bright pulse of right-circular light that he would be sure to detect it. Fortunately this attack succeeds so rarely (roughly with probability $\mu^3/32$) that it is a less serious threat than simple 2-photon beamsplitting.

In our setting, owing to the high transparency of the optical channel, only the first kind of beamsplitting attack is relevant, which, under the conservative assumptions explained above, leaks each successful bit to Eve with probability μ . If the quantum transmission consists of N successful pulses, Alice and Bob can therefore conservatively estimate that Eve has learned less than

$$N\mu + 5\sqrt{N\mu(1-\mu)}$$

bits through beamsplitting. The second term in this expression is, as before, a 5 standard-deviation allowance for Eve's having had better than average luck in her beam-splitting attempts. Our apparatus simulates beamsplitting without additional hardware simply by having the software Alice disclose directly to the software Eve the correct polarization of each successful pulse with probability μ .

4.3 Estimating Eve's Information

The expected fraction of Alice's string leaked to Eve through both kinds of eavesdropping is conservatively bounded above by

$$\rho = \mu + (4/\sqrt{2})p ,$$

where μ is the pulse intensity at the upstream end of the channel and p is the bit error rate. This formula comes from the fact that Eve can learn a fraction roughly μ of the bits through beamsplitting, and a number of effective bits no greater than

a fraction $(4/\sqrt{2})p$ from intercept/resend, as discussed above and in the Appendix. Because of possible correlation between these two kinds of information⁴, the total information gained by Eve will typically be less than the sum of these two terms. Because undiscovered errors can sometimes disappear without notice during the reconciliation protocol when bits are discarded to avoid leaking information to Eve, it is necessary for Alice and Bob to estimate the number of such undiscovered errors in order to estimate the value of p , and thus that of ρ . This can be achieved by a rather simple interpolation.

This estimate on ρ assumes that Eve has the superior technology required to delay measurement until after announcement of the correct bases (if she did not, the first term in the above formula for ρ would be decreased from μ to $\mu/\sqrt{2}$), that she has perfect photodetection, and perhaps most importantly that all transmission errors result from intercept/resend eavesdropping, rather than noise sources beyond Eve's control. These assumptions will in most cases be excessively conservative: eg in our case, many of the bit errors can be confidently attributed to causes other than eavesdropping.

If N is the number of successful pulses, Alice and Bob can conservatively estimate the number of bits leaked to Eve as

$$\ell = N\rho + 5\sqrt{N(\mu(1-\mu) + (4 + 2\sqrt{2})p)},$$

where the first term represents Eve's expected information and the second is a 5σ allowance for sampling errors. The two terms under the square root represent contributions to the variance of Eve's information from variances in, respectively, the number of split pulses at fixed μ and the number of effective bits leaked to Eve through intercept/resend. We take the square root of the sums rather than the sum of the square roots (ie we add the variances rather than the standard deviations) because the probability of success of beamsplitting is independent from the probability of success of intercept/resend, provided we consider only the successful pulses.

5 Sample Data from the Apparatus

Here we give examples of data actually transmitted through the quantum channel, the subsequent public discussion, and the size of the shared secret key ultimately distilled. The first batch of data is from a run in which there was in fact no eavesdropping, but the eavesdropper's potential information was nevertheless conservatively estimated as described above from the estimated error rate and known pulse intensity. The second batch of data illustrates the ability to distill a small amount of shared secret key from a run with significant amounts of both kinds of eavesdropping.

⁴ It would be hard for Eve to prevent learning some bits *twice*, ie through intercept/resending of a successfully split pulse.

5.1 A Run Without Eavesdropping

Here is some of the raw data obtained from data exchanged over the quantum channel on 27 February, 1991.

```
Alice 0001011110 1110110001 1000001000 0110011111 1110111110 0101011000 0111111101 0000111111 1111010101 0000101101
0011011101 0100001001 1111001101 1000000010 1011101010 0011011111 1010001001 0000011010 1101010000 1011001100
0101010010 0101010011 0000000010 0001100111 1011011011 0101100011 1100111011 1100111101 0000010110 0100000010
1000011110 0010110000 1000111111 1111101011 1010001110 1101000111 1010100100 0001010100 1010000010 0101010101
:
1101001110 1100000111 0001111000 1010010010 1101000100 0011111010 1110101100 0000111011 1110101001 1010011011

Bob 0001011100 1110110001 1000001000 0110011111 1110111110 0101011000 0111111101 0000111111 1111010101 0000101101
0011011101 0110001001 1111001101 1000000010 1011101010 0011011111 1010001001 0000011010 1101010000 1011001100
0101010010 0101010011 0010000010 0001100111 1011011011 0101100011 1100111011 1100111101 1000010110 0110000010
1000011110 0010110000 1000111111 1111101011 1010001110 1101000111 1010100100 0001011100 1000000010 0101011101
:
1101001110 1100000111 0001111000 1010010010 1101000100 0011111010 1110101100 0000111011 1110101001 1010011011
```

In this first example, out of about 715,000 pulses of intensity $\mu = 0.12$ sent by Alice, 2000 were received in the correct basis by Bob. This quantum transmission took about ten minutes of real time. Bob's string contained 79 errors, an error frequency of 3.95%.

A random permutation and block parity comparison was performed with block size 14, reducing the string length to 1678 bits with 29 remaining errors.

A second random permutation and block parity comparison was performed with block size 20, reducing the string length to 1483 bits with 4 remaining errors.

A third random permutation and block parity comparison was performed with block size 28, reducing the string length to 1420 bits with 2 remaining errors.

Random subset parity comparison was then begun, revealing an error on the first attempt. Removal of the error reduced the string length to 1410 bits, with one remaining error.

Random subset parity comparison was resumed, revealing an error on the second attempt. Removal of the error reduced the string length to 1399 bits, with no remaining errors.

Twenty random subset parities were then computed and found to agree, confirming to Bob and Alice that with high probability their remaining strings, now 1379 bits long, were identical.

A total of 76 errors were discovered during reconciliation. Using a simple interpolation formula, Bob and Alice estimated that about 3.6 errors had been eliminated without noticing, when bits were discarded in order to avoid leaking information to Eve during reconciliation. Hence, the original error rate was estimated to be 3.98%. Potential information leakage to Eve (and standard deviation) was therefore estimated to be equivalent to 466 (± 27.5) deterministic bits, comprising 226 bits from intercept/resend and 240 bits from beamsplitting, based on pulse intensity $\mu = 0.12$.

Therefore, allowing 159 bits ($5\sigma + 21$) excess compression for safety, it was decided to compress the string 625 bits by random subset hashing, leaving 754 bits of shared secret key distilled from 2000 original bits. Eve's expected information on this key was less than 10^{-6} bit, based on the privacy amplification theorem and the probability of a 5σ statistical deviation. In reality, of course, Eve has no information at all on this key since she did not eavesdrop on the quantum transmission, and the public reconciliation protocol is designed not to increase her information subsequently.

5.2 A Run With Substantial Eavesdropping

Substantial eavesdropping was attempted in the second example. Out of another approximately 715,000 light pulses of intensity $\mu = 0.12$ sent by Alice, 2000 were received in the correct basis by Bob. Through attempting to beamsplit all the pulses, and intercept/resending one eighth of them, the simulated Eve learned 336 individual bits of Alice's data, while increasing Bob's error frequency to 8.00% (160 errors). Eve read her selected pulses in random canonical bases, using a separate portion of the random bit diskette to make her decisions.

A random permutation and block parity comparison was performed with block size 7, reducing the string length to 1424 bits and leaving 49 errors⁵. Eve's information about the remaining string was still no more than 336 bits, and included knowledge of 242 individual bits.

A second random permutation and block parity comparison was performed with block size 10, reducing the string length to 1153 bits and leaving 11 errors. Eve's information about the remaining string was still no more than 336 bits, and included knowledge of 192 individual bits.

A third random permutation and block parity comparison was performed with block size 14, reducing the string length to 1027 bits and leaving no errors. Eve's information about the remaining string was still no more than 336 bits, and included knowledge of 167 individual bits.

Twenty consecutive successful random subset parity comparisons with no failures convinced Alice and Bob that their strings, now consisting of 1007 bits, were very probably identical. Eve's information about the remaining string was still no more than 336 bits, and included knowledge of 164 individual bits.

From the 148 errors found during reconciliation, Alice and Bob estimated that about 14 errors had been eliminated without noticing. Hence, the original error rate was estimated to be 8.10%. Potential information leakage to Eve (and standard deviation) was therefore estimated to be equivalent to 699 (± 36.3) deterministic bits, comprising 459 bits from intercept/resend and 240 bits from beamsplitting, based on pulse intensity $\mu = 0.12$.

⁵ They start with a block size different from that of the previous example because of reason (1) given at the beginning of Subsection 5.3.

Therefore, allowing 203 bits ($5\sigma + 21$) excess compression for safety, it was decided to compress the string 902 bits by random subset hashing, leaving 105 bits of shared secret key distilled from 2000 original bits. The 902-bit compressive hashing used to obtain the key was based on very conservative estimates of what Eve might know. Since she actually knew only 336 bits, this compression sufficed to reduce her expected information about the key to a ridiculously low level of $2^{-(902-336)}/\ln 2$, or about 6×10^{-171} bit.

5.3 Additional Remarks

The 2000-bit batch size used above for illustrative purposes is less than optimal. In production use, a larger batch size should be used for three reasons: 1) It would allow the users, by preliminary sampling, to get a good estimate of the bit error rate and so optimize the choice of block sizes used in the reconciliation stage; 2) by reducing the statistical uncertainty in estimating Eve's possible information, it would reduce the proportional amount of compression needed in the privacy amplification stage to assure a given level of security; and 3) if authentication of the public channel messages is necessary, the amount of key Alice and Bob must use up for this purpose is independent of batch size, hence the key expands by a larger factor the larger the batch size.

Let us finally recall that the reconciliation protocol described in this paper and used in our two examples is not the best possible. We ran our currently best protocol on the same data with the following results. The data in which 79 errors occurred in 2000 bits was reconciled at the cost of disclosing 530 bits of information (instead of 601 bits with the other protocol). Taking account of the subsequent 20 random subset parities to confirm success of the reconciliation and of the compression by 622 bits⁶ needed to eradicate the eavesdropper's information, the final secret key, the same for Alice and Bob, was 828 bits in length (instead of 754). The data in which 160 errors occurred in 2000 bits was reconciled at the cost of disclosing 868 bits of information (instead of 973). Taking account of the subsequent 20 random subset parities to confirm success of the reconciliation and of the compression by 895 bits needed to eradicate the eavesdropper's information, the final secret key, the same for Alice and Bob, was 217 bits in length (instead of 105). We do not describe this improved protocol here because it is still experimental and because it may well be that we shall discover yet a better protocol that bears no or little resemblance to our currently best protocol. The final protocol will be discussed in a subsequent paper.

⁶ We compress by 622 bits rather than 625 bits because one advantage of the new protocol is that it determines the exact number of errors, hence the potential leakage from intercept/resend is estimated more accurately.

Appendix :

Intercept/Resend and Privacy Amplification

The first part of this Appendix discusses in more detail the tradeoff between information gained and disturbance caused when Eve intercepts a pulse from Alice and resends it to Bob. The second part generalizes the privacy amplification techniques of [12] to show how Alice and Bob can nearly eradicate Eve's information. The generalization is necessary because [12] deals with deterministic information, whereas the information Eve gains through intercept/resend is generally of a probabilistic nature. For simplicity, we shall assume throughout this Appendix that Alice's pulses are pure single-photon states. This is justified because the threat created by multi-photon pulses is taken care of in the analysis of the beamsplitting attack (Subsection 4.2).

A.1 Error Rates for Eve and Bob in Intercept/Resend

Recall that the light pulses comprising the quantum transmission are prepared in a random sequence of the four canonical polarizations. In the intercept/resend attack, Eve intercepts selected pulses and measures them in bases her choosing, then she fabricates and sends to Bob other pulses in their place. Subsequently, too late to influence her choices of measurement, Eve is told whether the original pulses were from the rectilinear or circular bases. Eve would of course like to learn the polarization of the intercepted pulses exactly, and fabricate exact copies of them, but the uncertainty principle prevents her from doing so. Here we sketch a proof that the intercept/resend attack tells Eve the value of Alice's bits with probability at most $(2 + \sqrt{2})/4 \approx 85\%$, while inducing an error with probability at least 25% for each fabricated pulse that is later successfully measured by Bob in the correct original basis.

Before proceeding it should be noted that the *kind* of information Eve receives about Alice's bit depends on the basis in which Eve makes her measurement. If she uses either canonical basis, her information will be deterministic, since the subsequent public announcement of the correct basis will tell her either that she has learned Alice's bit correctly (having measured it the correct basis), or that she has spoiled the bit and knows nothing about it. On the other hand, if Eve uses an intermediate basis halfway between rectilinear and circular (the so-called Breidbart basis [7], to be described below), her information remains probabilistic, consisting of a measurement result that agrees with Alice's bit with probability $\wp = (2 + \sqrt{2})/4 \approx 85\%$ regardless of the announced basis. Even though the Breidbart measurement provides less Shannon information, $1 + \wp \log_2 \wp + (1 - \wp) \log_2(1 - \wp) \approx 0.399$ bit, than the $\frac{1}{2}$ expected bit of deterministic information provided by a measurement in rectilinear or circular bases, it is conceivable (see second part of the Appendix) that this probabilistic information requires Alice and Bob to throw away more information from the raw quantum transmission during privacy amplification than if Eve had used canonical bases.

The remainder of this part of the appendix justifies the lower bounds 25% and $\approx 15\%$ claimed above for Bob's and Eve's respective error probabilities on those successful pulses subjected to intercept/resent attack. It may be omitted by those uninterested in physical details of polarization measurements.

An arbitrary polarization state may be described by giving a point $Q = (X, Y, Z)$ on the unit sphere $X^2 + Y^2 + Z^2 = 1$, called in this context the *Poincaré sphere*. The parameters X , Y , and Z , represent respectively the rectilinear, diagonal, and circular components of state Q , and they can be determined with arbitrary accuracy for a bright beam of light consisting of many photons. However, the best that can be done to measure the polarization of a single photon is to cause it to interact with a measuring apparatus that forces it to choose between two states of a *basis*, ie a pair of states characteristic of the measuring apparatus, and represented by a pair of diametrically opposite points $\{P, -P\}$ on the Poincaré sphere. Upon encountering the measuring apparatus, a single photon in general behaves probabilistically: if a photon in state Q is measured in basis $\{P, -P\}$, it behaves like state P with probability $\cos^2(\alpha/2) = (1 + \cos \alpha)/2$, where α is the angle between points P and Q as seen from the center of the Poincaré sphere, and like state $-P$ with the complementary probability $\sin^2(\alpha/2) = (1 - \cos \alpha)/2$. Thus we see that the photon behaves deterministically precisely if $Q = \pm P$ because in this case $\alpha/2$ is either 0° or 90° . Subsequent measurements yield no further information about the original state Q , because the measuring apparatus necessarily either destroys the photon or forces it into one of the basis states $\pm P$ characteristic of the most recent measurement.

Two bases are *conjugate* [27] if the corresponding pairs of points are 90 degrees apart on the Poincaré sphere. The three standard mutually conjugate bases (rectilinear, diagonal, and circular) are conventionally identified with the intersections of the sphere with the X , Y , and Z axes respectively. Thus, we may think of polarization state $(1, 0, 0)$ as horizontal, $(-1, 0, 0)$ as vertical, $(0, 1, 0)$ as 45° , $(0, -1, 0)$ as 135° , $(0, 0, 1)$ as left-circular, and $(0, 0, -1)$ as right-circular. A state from one basis will behave randomly and equiprobably if measured in another basis conjugate to the first because $\cos^2 45^\circ = \frac{1}{2}$.

Suppose Alice sends a photon from the rectilinear basis $(\pm 1, 0, 0)$, which is intercepted and measured by Eve in an arbitrary basis $\{P, -P\}$, where $P = (X, Y, Z)$. Without loss of generality we assume X to be positive⁷. The measurement is counted as correct if it yields P when in fact the incoming photon was in polarization state $(1, 0, 0)$ or if it yields $-P$ when in fact the incoming photon was in polarization state $(-1, 0, 0)$; otherwise, the measurement is counted as an error. From the $\sin^2(\alpha/2)$ law, it is clear that Eve's error probability is $(1 - X)/2$ for either type of rectilinear photon. Suppose Eve next resends the photon in another basis $\{P', -P'\}$, according to the result $\{P, -P\}$ obtained in the first measurement, and finally that this photon is

⁷ If X is negative and the measured photon is rectilinear, Eve will obtain the correct result with probability less than $\frac{1}{2}$. However, she will discover this when the correct bases are announced on the public channel, and she can then negate her reading in order to obtain the correct bit with probability better than $\frac{1}{2}$. This is very much like using a binary symmetric channel with known error probability greater than $\frac{1}{2}$.

measured by Bob in the original rectilinear basis. Further application of the $\sin^2(\alpha/2)$ law shows that Bob's resulting error probability is $\frac{1}{4}[(1-X)(1+X')+(1-X')(1+X)]$. Similarly, for circular photons $(0, 0, \pm 1)$ sent by Alice and eventually received in the circular basis by Bob, the error probabilities induced by the same intercept/resend attack would be $(1-Z)/2$, and $\frac{1}{4}[(1-Z)(1+Z')+(1-Z')(1+Z)]$, for Eve and Bob respectively. The mean error probabilities averaged over both sending bases for Alice are thus

$$\frac{1}{4}[(1-X)+(1-Z)]$$

for Eve, and

$$\frac{1}{8}[(1-X)(1+X')+(1-X')(1+X)+(1-Z)(1+Z')+(1-Z')(1+Z)]$$

for Bob. From these formulas it can be verified that

- Eve's error probability, averaged over Alice's two sending bases, attains its minimum value when $X+Z$ is maximized, subject to $X^2+Y^2+Z^2=1$, which is clearly when $X=Z=\sqrt{2}/2$ and $Y=0$. In that case, Eve's error probability is $(2-\sqrt{2})/4 \approx 15\%$. This happens when Eve performs her measurement in a basis midway between the rectilinear and circular bases, henceforth called the Breidbart basis [7].
- Bob's error probability attains its minimum value of $\frac{1}{4}$ when Eve uses the same basis to resend as she used to intercept (ie $P'=P$), and this minimum value is achieved for any basis in the XZ plane. In particular Bob's error probability is $\frac{1}{4}$ when Eve intercept/resends in either the rectilinear, the circular, or the Breidbart basis. By contrast, Bob's error probability is necessarily larger than $\frac{1}{4}$ if $P' \neq P$ or if $Y \neq 0$, which means that Eve would create more errors than necessary if she does not resend in the interception basis or if she intercepts in a basis that contains a diagonal component. In particular, Bob's error will be exactly $\frac{1}{2}$ if Eve intercepts in the diagonal basis $(0, \pm 1, 0)$, regardless of her resending basis.
- Before public announcement of Alice's basis, Eve's *a priori* error probability on any given bit is $\frac{1}{4}[(1-X)+(1-Z)]$ as noted above. When the basis is announced, Eve will learn that her error probability for that bit was either $(1-X)/2$ or $(1-Z)/2$, according to whether the announced basis was rectilinear or circular. Her *a priori* information is maximized (at ≈ 0.399 bit) by intercept/resending in the Breidbart basis, while her expected *a posteriori* information is maximized (at $\frac{1}{2}$ bit) by intercept/resending in the rectilinear or circular basis.

Although we have implied above that any measurement by Eve on one of Alice's photons will induce an error with probability at least $\frac{1}{4}$ if the photon (or Eve's fabricated pulse) is later measured by Bob, this is only true if Eve's measuring apparatus

actually interacts with the photon. Eve could have a “measuring apparatus” that simply let the photon pass undisturbed, while also not telling her anything about it. Between these two extremes, Eve might have an apparatus that, whenever she attempted to use it, measured the photon with probability $s \leq 1$, otherwise letting it pass. Clearly such a sometimes-measurement yields no more than $s/2$ expected bit about the photon’s polarization, while inducing an error with probability at least $s/4$. Such a sometimes-measuring apparatus grants Eve no power she does not already have, by virtue of her ability to decide probabilistically whether or not to make an ordinary measurement.

A.2 Privacy Amplification Against Probabilistic Information

Assume that Eve performs intercept/resend on k successful pulses. If she uses canonical bases, she expects to learn the values of $k/2$ physical bits in Alice’s string, whereas if she uses any other basis she will learn less, and her information will be of a probabilistic nature. At first sight, Eve’s optimal strategy would thus appear to be to intercept/resend all the selected pulses in rectilinear or circular bases, a choice that also allows her to achieve minimal disturbance, $k/4$ expected errors, in the transmission reaching Bob. This reasoning would be valid if Eve wanted to maximize her information *on the raw quantum transmission*, but of course she wants to maximize her information on the final string shared between Alice and Bob after reconciliation *and privacy amplification*. The problem is that k bits known by Eve with probability roughly 85% each, even though they hold less information, turn out to be *more resistant* to privacy amplification than $k/2$ bits known with certainty.

Recall that the main privacy amplification theorem [12] says that if Eve knows ℓ *deterministic* bits of information about an n -bit string x , and if a suitable hash function $h : \{0, 1\}^n \rightarrow \{0, 1\}^{n-\ell-s}$ is then publicly but randomly chosen and applied to x , Eve’s expected information on the value of the hash function $h(x)$ will be less than $2^{-s}/\ln 2$ bit, where $s > 0$ is an arbitrary safety parameter. This theorem is not directly applicable to the situation in which Eve intercepts pulses in the Breidbart basis because in that case her information is probabilistic, consisting of knowledge of k bits of x with 85% probability. The difference between the two situations is that, in the case of deterministic information, Eve has a set of candidates, all equally plausible, for what x might be; whereas here her knowledge is characterized by a nonuniform probability distribution $p(x)$ in which candidates disagreeing with some of Eve’s Breidbart measurements receive correspondingly less weight, but are not excluded outright.

In [9] it is shown that such a nonuniform distribution of candidates resists privacy amplification no better than a set of $1/\sum_x p^2(x)$ equally weighted candidates, and therefore that k bits known with 85% probability, which supply $0.399 \times k$ bits of Shannon information, may resist privacy amplification somewhat better than $0.399 \times k$ bits of deterministic information, but can do so no better than

$(1 + \log_2(0.85^2 + 0.15^2)) \times k \approx 0.585 \times k$ deterministic bits⁸ Therefore, regardless of which basis Eve uses to intercept and resend, she cannot learn more than the privacy-amplification equivalent of 0.585 bits of deterministic information per successfully intercepted bit.

In the present paper, rather than justifying and using the above expressions for the true privacy-amplification equivalent of probabilistic information, we will derive and use a somewhat cruder upper bound of $k/\sqrt{2} = 0.7071 \times k$ bits on the amount of deterministic information required by Eve to simulate the effect of knowing k of Alice's bits with 85% probability. This approach involves introducing a new actor, whom we call Big Brother. We assume that Big Brother is given a fair coin, a biased coin whose bias is under his control, immediate access to all the bases used by Alice in forming her quantum transmission, and the ability to look at any of Alice's bits and learn its value reliably if he chooses to do so. Each time Eve decides to measure a pulse, Big Brother steps in and supplies her with a simulated reading (described below) that is statistically indistinguishable from the actual reading she would have obtained without his presence. The key point is that Big Brother does not always need to look at Alice's bit in order to supply Eve: sometimes he can get away with tossing a coin instead.

Consider first the simple case in which all of Eve's measurements are in the Breidbart basis. Each time that Eve selects a pulse for interception, Big Brother flips a biased coin. With probability $1/\sqrt{2}$, he looks at Alice's bit and gives it to Eve. Otherwise, he flips a fair coin and gives its outcome to Eve. Clearly, Eve gets Alice's bit with probability $(1/\sqrt{2}) \times 1 + (1 - 1/\sqrt{2}) \times \frac{1}{2} = (2 + \sqrt{2})/4$, exactly as she should have by intercepting in the Breidbart basis.

More generally, consider the case in which Eve measures Alice's photon in an arbitrary basis, having components X , Y , and Z , respectively of rectilinear, diagonal, and circular polarization, as explained in the preceding subsection. Using the formulas of that subsection, it is clear that Big Brother will correctly simulate Eve's data if he uses his prior knowledge of Alice's basis to decide between two courses of action.

- If Alice's basis was rectilinear, Big Brother looks at Alice's bit with probability X and feeds it to Eve; otherwise he feeds her a fair coin toss. In this case, Eve gets Alice's bit with probability $X \times 1 + (1 - X) \times \frac{1}{2} = (1 + X)/2$, exactly as she should.
- If Alice's basis was circular, Big Brother looks at Alice's bit with probability Z and feeds it to Eve; otherwise he feeds her a fair coin toss. In this case, Eve

⁸ In more detail, it is shown in [9], by arguments similar to those used in the context of quasi-perfect pseudorandom number generation by Impagliazzo and Zuckerman [21], that the property of a distribution most simply affected by privacy amplification is its Renyi or collision entropy $-\log_2 \sum_x p^2(x)$ which is less than the Shannon entropy except for uniform distributions where the two are equal. Hashing an input with Renyi entropy r down to an output of size less than r bits is necessary and sufficient to make the output's Renyi entropy near-maximal, which in turn forces the Shannon entropy to be near-maximal.

gets Alice's bit with probability $Z \times 1 + (1 - Z) \times \frac{1}{2} = (1 + Z)/2$, exactly as she should.

Clearly, Big Brother has to look at Alice's bit with probability $(X + Z)/2$. Subject to the constraint that $X^2 + Z^2 \leq 1$, this probability is maximized at $X = Z = \sqrt{2}/2$, in which case it takes the value $1/\sqrt{2}$. Thus we see that Eve forces Big Brother to spy hardest on Alice's bits precisely when she chooses to intercept in the Breidbart basis.

To complete the reasoning we note that, after all bases have been disclosed, Big Brother's expected information on Alice's bit string consists of $\ell \leq k/\sqrt{2}$ physical bits of Alice's data, which may be viewed as the value of an ℓ -bit deterministic function F of Alice's bits, where the function F is probabilistically chosen by Big Brother, who knows her bases but not her bits. Therefore Big Brother's information can be obliterated almost completely by privacy amplification. On the other hand, Eve knows no more about Alice's bits than what Big Brother has told her, which is generally less than all he knows. Therefore, Eve's information on Alice's string will also be almost obliterated by the same privacy amplification.

Of course, Alice and Bob need to estimate the number ℓ of bits obtained by Big Brother in order to apply privacy amplification. Because each pulse fabricated by Eve has at least a 25% chance of creating an error, and because Big Brother obtains a bit of Alice's string with probability no better than $1/\sqrt{2}$ when Eve attempts intercept/resend, it is clear that ℓ can be estimated as $(4/\sqrt{2})t$, where t is the number of errors between Alice's and Bob's data. However, a complication arises because Alice and Bob need to estimate not only the expected value of ℓ but also its standard deviation, in order to apply privacy amplification conservatively as if Big Brother had obtained 5 standard-deviation more bits than expected.

In order to calculate this standard deviation [22], let k denote the (unknown) number of pulses subjected to intercept/resend, let t denote the (observed)⁹ number of errors, and let ℓ denote the (unknown) number of bits leaked to Big Brother under the assumption that Eve used the optimal Breidbart basis for all her intercepts. It is clear that t is a Binomial($k, \frac{1}{4}$), whereas ℓ is a Binomial($k, 1/\sqrt{2}$). Let us now estimate ℓ by $\hat{\ell} = (4/\sqrt{2})t$. The variance of ℓ is $k(1/\sqrt{2})(1 - 1/\sqrt{2})$, which can be estimated as $4t(1/\sqrt{2})(1 - 1/\sqrt{2})$. The variance of our estimator $\hat{\ell}$ is $3k/2$, which can be estimated as $6t$ (it is $(4/\sqrt{2})^2$ times the variance of t). What we need, however, is the variance of the error in our estimate, ie the variance of $\ell - \hat{\ell}$. A calculation involving the central limit theorem shows that if k is large enough (about $k = 30$ suffices for all practical purposes), the distribution of $\ell - \hat{\ell}$ is nearly a normal of mean 0 and variance $k \times (2 + \sqrt{2})/2$. It follows that this variance can be estimated by $t \times (4 + 2\sqrt{2})$, and thus the standard deviation can be estimated as $\sqrt{(4 + 2\sqrt{2})t}$.

⁹ Using the reconciliation protocol described in this paper, t is estimated by interpolation rather than being actually observed. To be exact, we should have taken account also of that estimate's variance. We did not worry about this additional complication because our currently best reconciliation protocol (mentioned but not described in this paper) obtains the exact value for t .

A similar analysis can be made about Eve’s benefit from the beamsplitting attack if she does not have the superior technology required to store pulses for an arbitrarily long time before measuring them.

NOTE ADDED IN PROOF: Further work with Claude Crépeau [8] has shown that Eve’s optimal intercept/resend strategy, if she wishes to maximize her expected Shannon information on the final key string shared between Alice and Bob, is to measure her selected pulses in canonical (rectilinear or circular) bases. The estimate in Section 4.3 that as much as $1/\sqrt{2}$ bit of compression might be needed during privacy amplification to combat each intercept/resend was therefore overly conservative: 1/2 bit compression per intercept/resend would have sufficed. Accordingly, the privacy amplification used on the experimental data in Section 5 actually leaves Eve with far less than 10^{-6} bits expected information on the final key. It is interesting to note that if, instead of wishing to maximize her information on the key, Eve wishes to maximize her (very small) chance of guessing the entire key correctly, her optimal strategy is to conduct all her measurements in the intermediate Breidbart basis.

Acknowledgements

We are grateful to Grégoire Blet, André Chartier, Claude Crépeau, David Deutsch, Aviezri S. Fraenkel, N. David Mermin, Said Nsiri, and Stephen Wiesner for many helpful discussions. We are particularly indebted to Manuel Blum for his never-failing enthusiasm about quantum cryptography, to Joshua Rothenberg, Rodney Hodgson, and Leonard Mandel for advice on optical practice and theory, to Ueli Maurer for pointing out a conceptual error in a previous version of this paper (which tacitly assumed that Eve’s information was deterministic), and to Christian Léger for figuring out the standard deviation of the number of bits leaked to Big Brother. We are also grateful to an anonymous referee for several relevant suggestions, and to Ivan Damgård for his exemplary efficiency as guest editor of this *Journal of Cryptology*’s special issue after EUROCRYPT ’90.

References

- [1] Bengio, S., G. Brassard, Y. Desmedt, C. Goutier and J.-J. Quisquater, “Secure implementation of identification systems”, *Journal of Cryptology*, Vol. 4, no. 3, 1991.
- [2] Bennett, C. H. and G. Brassard, “An update on quantum cryptography”, *Advances in Cryptology: Proceedings of Crypto ’84*, August 1984, Springer-Verlag, pp. 475–480.
- [3] Bennett, C. H. and G. Brassard, “Quantum cryptography: Public key distribution and coin tossing”, *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, Bangalore, India, December 1984, pp. 175–179.
- [4] Bennett, C. H. and G. Brassard, “Quantum public key distribution system”, *IBM Technical Disclosure Bulletin*, Vol. 28, 1985, pp. 3153–3163.

- [5] Bennett, C.H. and G. Brassard, “The dawn of a new era for quantum cryptography: The experimental prototype is working!”, *Sigact News*, Vol. 20, no. 4, Fall 1989, pp. 78 – 82.
- [6] Bennett, C.H., G. Brassard and S. Breidbart, “Quantum cryptography II: How to re-use a one-time pad safely even if $\mathcal{P} = \mathcal{NP}$ ”, unpublished manuscript available from the authors, November 1982.
- [7] Bennett, C.H., G. Brassard, S. Breidbart and S. Wiesner, “Quantum cryptography, or unforgeable subway tokens”, *Advances in Cryptology: Proceedings of Crypto '82*, August 1982, Plenum Press, pp. 267–275.
- [8] Bennett, C.H., G. Brassard, C. Crépeau, and M.-H. Skubiszewska, “Practical quantum oblivious transfer”, *Advances in Cryptology — Crypto '91 Proceedings*, to appear.
- [9] Bennett, C.H., G. Brassard, C. Crépeau, and U.M. Maurer, “Privacy amplification against probabilistic information”, in preparation.
- [10] Bennett, C.H., G. Brassard and N.D. Mermin, “Quantum cryptography without Bell’s theorem”, *Physical Review Letters*, Vol. 68, no. 5, 1992, pp. 557–559.
- [11] Bennett, C.H., G. Brassard and J.-M. Robert, “How to reduce your enemy’s information”, *Advances in Cryptology — Crypto '85 Proceedings*, August 1985, Springer-Verlag, pp. 468–476.
- [12] Bennett, C.H., G. Brassard and J.-M. Robert, “Privacy amplification by public discussion”, *SIAM Journal on Computing*, Vol. 17, no. 2, April 1988, pp. 210–229.
- [13] Brassard, G., *Modern Cryptology: A Tutorial*, Lecture Notes in Computer Science, Vol. 325, Springer-Verlag, Heidelberg, 1988.
- [14] Brassard, G. and C. Crépeau, “Quantum bit commitment and coin tossing protocols”, *Advances in Cryptology — Crypto '90 Proceedings*, to appear.
- [15] Brickell, E.F. and A.M. Odlyzko, “Cryptanalysis: A survey of recent results”, *Proceedings of the IEEE*, Vol. 76, no. 5, May 1988, pp. 578–593.
- [16] Crépeau, C., “Correct and private reductions among oblivious transfers”, PhD Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, February 1990.
- [17] Crépeau, C. and J. Kilian, “Achieving oblivious transfer using weakened security assumptions”, *Proceedings of 29th IEEE Symposium on the Foundations of Computer Science*, White Plains, NY, October 1988, pp. 42–52.
- [18] Deutsch, D., “Quantum communication thwarts eavesdroppers”, *New Scientist*, December 9, 1989, pp. 25–26.
- [19] Ekert, A., “Quantum cryptography based on Bell’s theorem”, *Physical Review Letters*, Vol. 67, no. 6, 5 August 1991, pp. 661–663.
- [20] Gottlieb, A., “Conjugal secrets — The untappable quantum telephone”, *The Economist*, Vol. 311, no. 7599, 22 April 1989, p. 81.
- [21] Impagliazzo, R. and D. Zuckerman, “How to Recycle Random Bits”, *Proceedings of 30th IEEE Symposium on the Foundations of Computer Science*, Research Triangle Park, NC, October 1989, pp. 248–253.
- [22] Léger, C., personal communication.

- [23] Peterson, I., “Bits of uncertainty: Quantum security”, *Science News*, Vol. 137, 2 June 1990, pp. 342–343.
- [24] Robert, J.-M., “Détection et correction d’erreurs en cryptographie”, *Masters Thesis*, Département d’informatique et de recherche opérationnelle, Université de Montréal, Montréal (Québec), Canada, 1985.
- [25] Wallich, P., “Quantum cryptography”, *Scientific American*, Vol. 260, no. 5, May 1989, pp. 28–30.
- [26] Wegman, M. N. and J. L. Carter, “New hash functions and their use in authentication and set equality”, *Journal of Computer and System Sciences*, Vol. 22, 1981, pp. 265–279.
- [27] Wiesner, S., “Conjugate coding”, manuscript written *circa* 1970, unpublished until it appeared in *Sigact News*, Vol. 15, no. 1, 1983, pp. 78–88.